

APPLICATION
FOR
UNITED STATES LETTERS PATENT

APPLICANT NAME: Anthony R. Tuel

TITLE: METHOD, SYSTEM AND PROGRAM PRODUCT FOR
PROCESSING A TRANSACTION

DOCKET NO.: RSW920030189US1

INTERNATIONAL BUSINESS MACHINES CORPORATION

CERTIFICATE OF MAILING UNDER 37 CFR 1.10

I hereby certify that, on the date shown below, this correspondence is being deposited with the United States Postal Service in an envelope addressed to the Commissioner for Patents, Mail Stop: Patent Application, PO Box 1450, Alexandria, Virginia as "Express Mail Post Office to Addressee" Mailing Label No. EV225574345US

on October 30, 2003

Dorothea Rubbone

Name of person mailing paper

Dorothea Rubbone
Signature

10/30/2003
Date

METHOD, SYSTEM AND PROGRAM PRODUCT FOR PROCESSING A TRANSACTION

BACKGROUND OF THE INVENTION

1. TECHNICAL FIELD

[0001] The invention relates generally to processing transactions, and more specifically, to a solution for more efficiently processing and/or responding to a request for a transaction that requires multiple resources.

2. BACKGROUND ART

[0002] Many transactions require that operations be performed on multiple resource managers and/or multiple resources under the same resource manager. For example, a distributed XA transaction may be requested that requires entries to be added and/or changed in multiple databases. In these transactions, it is typically desired that either all the operations occur successfully or none of the operations occur at all. In order to ensure this, the operations can be performed using the XA protocol, which supports a two phase process. In the first phase, each resource is "prepared." Preparing a resource includes providing the operation to the resource manager, performing the operation on the resource, and obtaining and temporarily storing result(s) for the operation. The result(s) from the resource manager is then returned to a transaction manager that manages the entire transaction. When the returned results indicate that all operations were successful, each resource can be "committed" in the second phase. Committing a resource includes making the result(s) of the operation permanent in the resource. However, if the returned results indicate that one or more operations failed, then each resource

can instead be "rolled back" in the second phase. Rolling back a resource includes undoing the result(s) of the operation performed when the resource was prepared in the first phase as well as any operations performed on the resource prior to the preparation.

[0003] Typically, in processing a transaction that spans multiple resources, each resource is serially prepared and serially committed or rolled back. For example, when two resources are required, the first resource is prepared, and when a successful preparation response is received, the second resource is prepared. Assuming both resources are prepared successfully, the first resource is committed followed by committing the second resource. Once both resources have been successfully committed, the transaction manager provides a response to the requester.

[0004] Serially processing resources for the transaction can be inefficient. For example, one or more resources may have a separate resource manager that is concurrently performing actions for other systems. As a result, when the transaction manager requests that the resource manager perform an operation, there may be some delay before the operation is performed and a response is received. When numerous resources are required for a transaction, these delays can accumulate into a substantial delay in responding to the transaction request.

[0005] As a result, a need exists for an improved solution for processing transactions that require multiple resources. In particular, a need exists for a method, system and program product that more efficiently process a transaction by simultaneously preparing and/or committing resources before replying to the requester.

SUMMARY OF THE INVENTION

[0006] The invention provides an improved solution for processing transactions that require multiple resources. Specifically, under the present invention, a method, system and program product for processing transactions are provided in which the preparation and/or commitment of the resources are performed concurrently. In particular, after receiving a transaction request, preparation of all the resources for the transaction is requested without waiting to receive a preparation response for a previous resource. Each request can be made using a non-blocking function call or by using a unique thread. In any event, the resources are concurrently prepared. Should the resources require commitment or roll back, these operations can also be performed concurrently. In order to further improve processing, a transaction response can be sent to a requester before the resources are committed or rolled back.

[0007] A first aspect of the invention provides a method of processing a transaction that requires a plurality of resources, the method comprising: requesting preparation of a first resource; and requesting preparation of a second resource before receiving a preparation response from the first resource.

[0008] A second aspect of the invention provides a method of processing a transaction, the method comprising: concurrently preparing a plurality of resources for the transaction; waiting for a preparation response for each of the plurality of resources; and concurrently committing the plurality of resources.

[0009] A third aspect of the invention provides a system for processing a transaction that requires a plurality of resources, the system comprising: a reception system for receiving a transaction request from a requester; and a preparation system for concurrently preparing the plurality of resources for the transaction.

[0010] A fourth aspect of the invention provides a program product stored on a recordable medium for processing a transaction, which when executed comprises: program code for requesting preparation of a plurality of resources for the transaction; program code for simultaneously waiting for a preparation response for each of the plurality of resources; and program code for requesting at least one of: commitment and roll back of the plurality of resources.

[0011] The illustrative aspects of the present invention are designed to solve the problems herein described and other problems not discussed, which are discoverable by a skilled artisan.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] These and other features of this invention will be more readily understood from the following detailed description of the various aspects of the invention taken in conjunction with the accompanying drawings in which:

[0013] FIG. 1 shows an illustrative system for processing transactions according to one embodiment of the invention;

[0014] FIG. 2 shows an illustrative flow diagram for processing a transaction according to another embodiment of the invention; and

[0015] FIG. 3 shows an illustrative flow diagram for processing a transaction according to yet another embodiment of the invention.

[0016] It is noted that the drawings of the invention are not to scale. The drawings are intended to depict only typical aspects of the invention, and therefore should not be considered as limiting the scope of the invention. In the drawings, like numbering represents like elements between the drawings.

DETAILED DESCRIPTION OF THE INVENTION

[0017] As indicated above, the invention provides an improved solution for processing transactions that require multiple resources. Specifically, under the present invention, a method, system and program product for processing transactions are provided in which the preparation and/or commitment of the resources are performed concurrently. In particular, after receiving a transaction request, preparation of all the resources for the transaction is requested without waiting to receive a preparation response for a previous resource. Each request can be made using a non-blocking function call or by using a unique thread. In any event, the resources are concurrently prepared. Should the resources require commitment or roll back, these operations can also be performed concurrently. In order to further improve processing, a transaction response can be sent to a requester before the resources are committed or rolled back.

[0018] As used throughout this discussion, the term "transaction" is used to describe any type of desired action that requires one or more operations to be performed on one or more resources. The term "resources" is used to describe any type of input device, output device, storage device, etc. For example, a resource may comprise a database, a webserve queue, a Java Messaging Services (JMS) provider, a downstream server, etc. Further, it is understood that an "operation" on a resource could comprise multiple actions that are to be performed on the resource. For example, a transaction could comprise updating a table in one database, and updating two tables in another database. In this case, the transaction could be described as comprising an operation for one resource and an operation for a second resource. Still further, it is understood that various problems can arise to cause an operation on a resource to fail. For example, when the resource is a database, the operation could identify a table that does not exist, an entry that is not present in a table, etc. Further, the resource may not perform the operation within a

predetermined time period, insufficient memory may be available, the resource may be temporarily unavailable, etc.

[0019] Turning to the drawings, FIG. 1 shows an illustrative system 10 for processing transactions. In particular, a computer 12 can receive a transaction request from a requester 26 over a communications link 13A. In processing the transaction, computer 12 can request that one or more resource managers 40 perform an operation on one or more resources 42 over communications link 13B. To this extent, each communications link 13A-B can comprise a direct hardwired connection (e.g., serial port) or another type of network connection. In the latter case, the network can comprise an addressable connection in a client-server (or server-server) environment that may utilize any combination of wireline and/or wireless transmission methods. In this instance, computer 12, requester 26, and resource managers 40 may utilize conventional network connectivity, such as Token Ring, Ethernet, WiFi or other conventional communications standards. Further, the network can comprise any form of network, including the Internet, a wide area network (WAN), a local area network (LAN), a virtual private network (VPN), etc. Where, for example, requester 26 communicates with computer 12 via the Internet, connectivity could be provided by conventional TCP/IP sockets-based protocol, and requester 26 could utilize an Internet service provider to establish connectivity to computer 12.

[0020] As shown, computer 12 generally includes central processing unit (CPU) 14, memory 16, input/output (I/O) interface 18, bus 20, external I/O devices/resources 22, and a storage unit 24. CPU 14 may comprise a single processing unit, or be distributed across one or more processing units in one or more locations, e.g., on a client and server. Memory 16 may comprise any known type of data storage and/or transmission media, including magnetic media, optical media, random access memory (RAM), read-only memory (ROM), a data cache, a data object,

etc. Storage unit 24 may comprise any type of data storage for providing storage for information necessary to carry out the invention as described below. As such, storage unit 24 may include one or more storage devices, such as a magnetic disk drive or an optical disk drive. Moreover, similar to CPU 14, memory 16 and/or storage unit 24 may reside at a single physical location, comprising one or more types of data storage, or be distributed across a plurality of physical systems in various forms. Further, memory 16 and/or storage unit 24 can include data distributed across, for example, a LAN, WAN or a storage area network (SAN) (not shown).

[0021] I/O interface 18 may comprise any system for exchanging information to/from external devices. I/O devices 22 may comprise any known type of external device, including speakers, a CRT, LED screen, handheld device, keyboard, mouse, voice recognition system, speech output system, printer, monitor/display, facsimile, pager, etc. Bus 20 provides a communication link between each of the components in computer 12 and likewise may comprise any known type of transmission link, including electrical, optical, wireless, etc. In addition, although not shown, additional components, such as cache memory, communication systems, system software, etc., may be incorporated into computer 12.

[0022] Further, it is understood that computer 12 comprises any type of computing device capable of communicating with one or more other computing devices (e.g., requester 26), such as a server, a desktop computer, a laptop, a handheld device, a mobile phone, a pager, a personal data assistant, etc. Similarly, requester 26 can comprise any type of computing device. To this extent, requester 26 typically includes the same elements as shown in computer 12 (e.g., CPU, memory, I/O interface, etc.). These have not been separately shown and discussed for brevity. It is understood, however, that if computer 12 is a handheld device or the like, a display could be contained within computer 12, and not as an external I/O device 22 as shown. Still further, each

resource manager 40 can comprise a component in a system that includes the managed resource(s) 42. Each resource manager 40 can communicate with other computing devices, and can manage the processing of operations on one or more resources 42. To this extent, each resource manager 40 can manage the implementation of the two phase processing (e.g., preparation followed by commitment or roll back) on a resource 42.

[0023] Shown stored in memory 16 is a processing system 28 for processing transactions. Processing system 28 is shown including a reception system 30, a preparation system 32, a logging system 34, a commitment system 36, and a reply system 38. In general, reception system 30 can receive a transaction request for a transaction that requires multiple resources 42. Preparation system 32 can process the preparation phase of the transaction, logging system 34 can log the results of the preparation of resources 42, and commitment system 36 can process the commitment/roll back phase of the transaction. Reply system 38 can reply back to, for example, requester 26 with a transaction result.

[0024] As noted, reception system 30 can receive a transaction request from, for example, requester 26. Alternatively, it is understood that a user (not shown) could operate computer 12 and generate one or more transaction requests that are sent to processing system 28 for processing, and are received by reception system 30. In any event, when reception system 30 receives a transaction request, it can determine the number of resources 42 that are required for the transaction. For example, the transaction request may identify each required resource 42, or reception system 30 can analyze the transaction request to determine the required resource(s) 42. When a transaction request only requires a single resource 42, processing system 28 can process the transaction in a single phase, i.e., the result of the operation for the transaction is

permanently stored in resource 42 immediately. This can be performed without using preparation system 32 and commitment system 36.

[0025] However, when reception system 30 determines that the requested transaction requires multiple resources 42, it can forward the transaction request to preparation system 32 for processing the preparation phase of the transaction. To this extent, preparation system 32 can request the resource manager 40 of a resource 42 required for the transaction to prepare the corresponding resource 42. Further, preparation system 32 could prepare one or more of the required resources 42 without the use of a resource manager 40. In any event, the required resources 42 for the transaction are concurrently prepared.

[0026] In one embodiment, preparation system 32 comprises a single execution thread that requests preparation of each resource 42 without waiting to receive a preparation response from a previous resource 42. To this extent, each request could be made in a non-blocking manner. For example, preparation system 32 can use a non-blocking function call that requests resource manager 40 to prepare resource 42, but does not wait to receive a preparation response from the resource managers 40. Alternatively, a non-blocking standard query language (SQL) call or the like could be made to a resource 42 directly from the execution thread. Once all resources 42 are being prepared, preparation system 32 can wait for a preparation response to be received for each resource 42.

[0027] In another embodiment, preparation system 32 can use a unique resource thread for each resource 42. For example, preparation system 32 can comprise a main execution thread that starts a unique resource thread to process the preparation of each resource 42. To this extent, a resource thread can request that a resource manager 40 prepare the resource 42, and wait to receive the preparation response for the resource 42. Alternatively, a resource thread

could prepare the resource 42 by, for example, making a SQL call that blocks, and returns a preparation response once the resource has been prepared. In any event, once the resource 42 has been prepared, and a preparation response is available, the resource thread can return the preparation response to the main execution thread and terminate. Once all the resource threads have returned preparation responses, the main execution thread can continue processing the transaction.

[0028] For further processing of the transaction, preparation system 32 can provide all the preparation responses to logging system 34. Logging system 34 can log a preparation result for the transaction based on the preparation responses. In particular, logging system 34 can analyze the various preparation responses and determine a preparation result for the transaction. For example, a preparation response may indicate that the operation was successful or failed, and/or may return data for the operation. This information can be collected for the various resources 42 and compiled into a preparation result for the transaction.

[0029] After logging system 34 has generated a preparation result for the transaction, commitment system 36 can commit or roll back the various resources 42. In particular, if one or more of the operations on resources 42 failed during preparation, commitment system 36 can roll back all resources 42 to their state before the operations were performed. Alternatively, when all the operations on resources 42 are successfully prepared, commitment system 36 can commit the results of the operations on each resource 42, thereby making them permanent.

[0030] Similar to preparation system 32, commitment system 36 can concurrently commit or roll back resources 42. In particular, commitment system 36 can comprise a single execution thread that requests commitment of all resources 42 without waiting for a commitment response from any resource, and subsequently waits to receive all commitment responses. Alternatively,

commitment system 36 can comprise a main execution thread that creates or retrieves from a thread pool a unique resource thread for each resource 42. Each resource thread can request commitment of resource 42, wait to receive a commitment response for resource 42, return the commitment response to the main execution thread, and terminate. It is understood that commitment system 36 can also roll back resources 42 in an identical manner.

[0031] When preparation system 32 and commitment system 36 both use resource threads for preparing and committing/rolling back resources 42, it is understood that the same resource threads could be used for both operations. For example, after preparing a resource 42 and providing the preparation result to the main thread, a resource thread can wait to receive a command from the main execution thread to commit or roll back resource 42. In this case, the main execution thread can wait for all the preparation results, determine whether preparation of one or more resources 42 failed, and send the command to each resource thread to either commit or roll back the corresponding resource 42. Subsequently, the main execution thread can wait to receive a commitment (roll back) result for each resource 42 before continuing processing the transaction.

[0032] As noted previously, reply system 38 sends a reply to, for example, requester 26 based on a transaction result. The transaction result can be based on the preparation results and/or the commitment (roll back) results for resources 42. In particular, when the preparation of one or more resources 42 fails, the transaction result can indicate that preparation of the transaction failed. However, when all resources 42 are prepared successfully, the transaction result can comprise the preparation result. Further, the transaction result may indicate that one or more resources 42 were unable to be committed (rolled back) after a certain amount of time.

[0033] However, in many systems, it can be safely assumed by requester 26 that all resources 42 will be successfully committed or rolled back. For example, a resource 42 may only fail to be successfully committed when an heuristic exception occurs, e.g., it loses power or the like. In this case, resource manager 40 for resource 42 and/or processing system 28 will recognize this failure when power is restored to resource 42 and commit the results. Because of this, reply system 38 may reply to requester 26 before resources 42 have been committed to further decrease the amount of time to that requester 26 must wait for a reply.

[0034] FIG. 2 shows an illustrative flow diagram for processing a transaction that requires multiple resources using a single execution thread. As shown, a transaction request is received in step S1, and in step S2, the execution thread requests preparation of resources 42 (FIG. 1). In steps R1A-B, resources 42 are prepared concurrently by the corresponding resource managers 40 (FIG. 1), while in step S3, the execution thread waits for all the preparation responses. In step S4, the execution thread determines a preparation result for the transaction based on the preparation responses and logs the preparation result. As noted previously, in step S5, the execution thread can reply to requester 26 (FIG. 1) based on the preparation result before resources 42 are committed. In step S6, the execution thread requests that resources 42 be committed, and in steps R2A-B resources 42 are concurrently committed by resource managers 40, while in step S7, the execution thread waits to receive all commitment responses before continuing processing.

[0035] FIG. 3 shows another illustrative flow diagram for processing a transaction that requires multiple resources using multiple threads. As shown, in step M1, a main execution thread can receive a transaction request, and in step M2, the main execution thread can create a unique resource thread for each resource 42 (FIG. 1). In steps T1A-B, the resource threads can request

preparation of resources 42. In steps T2A-B, each resource thread can wait for a preparation response for the corresponding resource 42, while in step M3, the main execution thread can wait to receive all the preparation responses for resources 42 from the resource threads. In step M4, the main execution thread determines a preparation result for the transaction and logs the preparation result, while in steps T3A-B, each resource thread waits for a command from the main execution thread to either commit or roll back resources 42. Assuming the preparation result indicates that resources 42 were successfully prepared, the main execution thread can request that resources 42 be committed in step M5. In steps T4A-B, each resource thread can request that the corresponding resource 42 be committed. In steps T5A-B, each resource thread can wait for a commitment response for the resource 42, while in step M6, the main execution thread can wait to receive commitment responses for all resources 42. In step M7, the main execution thread can reply to the requester once all resources 42 have been successfully committed.

[0036] It is understood that both flow diagrams shown in FIGS. 2 and 3 are merely illustrative, and various alternatives are possible. For example, in FIG. 2, the reply can be sent after all resources 42 have been committed. Similarly, in FIG. 3, the reply can be sent before resources 42 are committed. Additionally, only the preparation phase or the commitment phase may be performed concurrently. Still further, the main execution thread could continue processing and allow one or more separate threads to wait for resources 42 to be successfully committed.

[0037] Further, while shown and discussed with reference to two resources, it is understood that a transaction can require any number of resources. Still further, it is understood that the present invention can be realized in hardware, software, or a combination of hardware and software. Any kind of computer/server system(s) - or other apparatus adapted for carrying out

the methods described herein - is suited. A typical combination of hardware and software could be a general-purpose computer system with a computer program that, when loaded and executed, carries out the respective methods described herein. Alternatively, a specific use computer (e.g., a finite state machine), containing specialized hardware for carrying out one or more of the functional tasks of the invention, could be utilized. The present invention can also be embedded in a computer program product, which comprises all the respective features enabling the implementation of the methods described herein, and which - when loaded in a computer system - is able to carry out these methods. Computer program, software program, program, or software, in the present context mean any expression, in any language, code or notation, of a set of instructions intended to cause a system having an information processing capability to perform a particular function either directly or after either or both of the following: (a) conversion to another language, code or notation; and/or (b) reproduction in a different material form.

[0038] The foregoing description of various aspects of the invention has been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed, and obviously, many modifications and variations are possible. Such modifications and variations that may be apparent to a person skilled in the art are intended to be included within the scope of the invention as defined by the accompanying claims.